



Terrain Classification for Enhanced Autonomous Systems

CH Bhavya Sai, V Sai Surya Teja, A Prudhvi, G Kushwanth, B Bikash

Department of Computer Science & Engineering

GMR Institute of Technology, Rajam, Vizianagaram

Vssteja2003@gmail.com

DOI: <https://doi.org/10.55248/gengpi.5.0324.0776>

ABSTRACT –

Terrain classification is a critical component in numerous applications, spanning robotics, autonomous vehicles, and military operations. It involves categorizing different terrains based on their physical attributes, such as texture, elevation, and surface composition. This categorization enables machines and systems to comprehend and adapt to various landscapes, facilitating informed decisions on navigation and environmental interaction. Achieving precise terrain classification relies on a variety of techniques, including deep learning algorithms, transfer learning methods, Auto Encoders, and Vision Transformers. These approaches leverage data from sensors like LiDAR, cameras, and radar to discern ground characteristics accurately. By distinguishing between categories such as flat surfaces, inclines, vegetation, water bodies, and obstacles, these systems bolster their navigation and decision-making capabilities. Accurate terrain classification is essential for enhancing autonomous systems' navigation and decision-making, particularly in path planning, obstacle avoidance, and situational awareness. For instance, autonomous vehicles rely on precise terrain classification for route planning and obstacle detection, while robots involved in disaster relief efforts depend on it for navigating safely through challenging environments. Ongoing research endeavors in terrain classification aim to enhance the resilience and efficiency of these systems across diverse environmental contexts. As technology progresses, there's an increasing demand for more dependable and effective terrain classification systems capable of operating in various environmental conditions.

Keywords— Terrain Classification, Autonomous Systems, Navigation, LiDAR, surface composition

1. INTRODUCTION

Terrain classification is pivotal across various domains, notably in robotics, autonomous vehicles, and military operations, where effective navigation relies on understanding diverse landscapes. Leveraging sensor data, camera data and acoustic information, terrain classification enables machines to discern terrain characteristics crucial for informed decision-making. This study focuses on leveraging deep learning techniques to accurately classify terrains, with potential applications in military applications and decision-making. This paper delves into terrain classification methodologies, utilizing transfer learning models and vision transformers augmented with preprocessing techniques. By harnessing deep learning algorithms and sensor data analysis, these models distinguish terrain features like slopes, vegetation, and obstacles, enhancing navigation and situational awareness for autonomous systems.

Notably, terrain classification emerges as a critical challenge, underscored by its inclusion in the Ministry of Defense's problem statement during the Smart India Hackathon 2023, highlighting its significance in addressing real-world needs. Additionally, the dataset used in this study is recommended by the Ministry of Defense, ensuring relevance and applicability to military applications.

2. BACKGROUND

[1] Yang et al. (2023) delve into the domain of terrain classification with a focus on leveraging deep learning techniques for automated classification using high-resolution Digital Elevation Model (DEM) data. By employing semantic segmentation models like Fully Convolutional Networks (FCN) and ResNet, they aim to accurately categorize landforms based on terrain factors such as slope and aspect. Their study showcases stable performance metrics with Precision Accuracy (PA) at 80% and Mean Intersection over Union (MIoU) at 68%, primarily utilizing DEM data. However, challenges arise in delineating boundaries between landform types accurately, especially in transition zones. The research hints at future explorations in incorporating auxiliary data to enhance classification accuracy and extending deep learning methodologies to classify genetic landform types alongside morphological ones.

[2] Bai et al. (2019) propose a vibration-based terrain classification system tailored for mobile robots. Their approach, based on Fast Fourier Transformation (FFT) and BP neural networks, achieves impressive accuracy rates, notably 98% across five terrain types. However, challenges persist, particularly in accurately classifying sand terrain, indicating the need for further refinement. The authors highlight the importance of sensor integration

and fusion to enhance classification accuracy and robustness, underlining the potential for future research in exploring deep learning techniques for analyzing sound and vibration data.

[3] Shi et al. (2020) delve into vibration-based robotic terrain classification, focusing on a modified Laplacian Support Vector Machine (SVM) approach. Their study addresses limitations of traditional SVMs with a semi-supervised learning framework, achieving significant improvements in classification accuracy. However, challenges remain regarding the robustness of the proposed method under varying conditions and the absence of comprehensive computational complexity analysis. Future research directions include exploring semi-supervised learning further, integrating terrain vision, and incorporating temporal dependencies for enhanced classification.

[4] Christie and Kottege (2016) present a real-time terrain classification system for legged robots based on acoustic data analysis. Their approach, utilizing acoustic recordings from various terrains, achieves high sensitivity rates. However, challenges arise in distinguishing between certain terrains, suggesting potential overlap in auditory interactions and feature space. The authors suggest future testing of the system on robotic platforms in diverse environments to evaluate its practical applicability and performance under varied conditions.

[5] Pingel et al. (2013) introduce an improved morphological filtering algorithm for terrain classification using LiDAR data. Their algorithm, SMRF, focuses on achieving high accuracy in ground filtering and terrain modeling for immersive virtual environments. While demonstrating promising results, the study acknowledges limitations in comparative analysis and suggests future improvements in addressing Type II errors to enhance visual quality in generated models.

[6] Li et al. (2020) propose a deep learning-based approach for landform classification, integrating data from Digital Elevation Models (DEMs) and remote sensing imagery. Their methodology outperforms traditional methods, showcasing higher accuracy and distinct landform boundaries. However, the study highlights dependencies on specific datasets for network training and ambiguity in landform boundaries as key limitations, suggesting future work in defining optimal boundaries for improved classification.

[7] Zhang et al. (2018) explore 3D Convolutional Neural Networks (CNN) for terrain classification using polarimetric SAR data. Their study demonstrates superior performance of 3D CNN over 2D CNN in terrain classification. Challenges include navigating complex cost functions and optimizing computational efficiency of 3D CNNs while maintaining robustness. Future research may focus on improving efficiency and robustness of 3D CNNs in terrain classification.

[8] Yu (2020) investigates remote sensing image terrain classification using an improved K-nearest neighbors (KNN) algorithm. By combining AlexNet feature extraction and KNN classification, the study achieves enhanced accuracy in image classification. Future trends may involve larger-scale data usage, deeper neural network architectures, and more efficient numerical optimization methods for remote sensing image classification.

[9] Zürn et al. (2020) propose a self-supervised visual terrain classification method for autonomous robots, aiming for improved clustering accuracy and reduced label noise. Their approach, utilizing triplet loss objectives and large-scale multimodal datasets, demonstrates promising clustering accuracies. Challenges include accurately predicting semantic classes in challenging visual environments and the requirement for manual labeling of data samples. Future research directions involve enhancing self-supervised learning for outdoor robotic environments.

[10] Abete et al. (2021) introduce a terrain identification method for humanoid robots using convolutional neural networks (CNN) and tactile perception through sensor data. Their approach achieves high accuracy in classifying various terrain types, outperforming existing methods for legged robots. Challenges include disruptions in classification with changes in walking velocities and variations in ground types, suggesting future exploration of recurrent neural network architectures and cost-effective sensor data reduction methods.

3. METHODOLOGY

Dataset Collection and Preparation:

The study commenced by accessing a dataset sourced from Kaggle, representing real-time terrain data compiled from diverse environments. This dataset encompassed various terrain classes such as grassy, marshy, rocky, and sandy terrains. To ensure comprehensive analysis, the dataset underwent meticulous processing, including division into an initial origin dataset and an augmented dataset. This division aimed to enhance the dataset's diversity and representativeness, enabling robust training and evaluation of terrain classification models. Description of the dataset is as follows.

Data Augmentation:

Data augmentation involves creating new training samples by applying various transformations to the existing dataset. Here's how the process works:

Rescaling: The pixel values of the images are rescaled to a range between 0 and 1. This normalization ensures uniformity in the input data.

Rotation: Each image is randomly rotated within a specified range, in this case, up to 20 degrees. This variation helps the model generalize better to different orientations of objects in real-world scenarios.

Horizontal and Vertical Shift: Random shifts are applied to the images horizontally and vertically. The shifts are proportional to the image size, with a range of 20% of the image width and height, respectively. This augmentation simulates variations in the position of objects within the images.

Horizontal Flipping: Images are randomly flipped horizontally. This transformation adds diversity to the dataset by presenting mirror images of the original samples.

Fill Mode: This parameter specifies how new pixels should be filled in when the image is shifted or rotated. The 'nearest' mode fills new pixels with the value of the nearest pixel in the original image.

Overall, this data augmentation process enriches the training dataset with a wider variety of samples, helping the model learn robust features and improve its performance when faced with unseen data during training.

Model Building:

Transfer learning models, such as VGG16, DenseNet-121, MobileNet, Inception V3, ResNet-101, DenseNet-201, and ResNet-50, are widely used in various computer vision tasks due to their pre-trained weights on large-scale datasets like ImageNet. VGG16 is characterized by its simplicity and consists of multiple convolutional and fully connected layers. DenseNet-121 employs densely connected blocks to promote feature reuse and parameter efficiency. MobileNet is designed for mobile and embedded applications, featuring lightweight depth-wise convolutions. Inception V3 utilizes inception modules with multiple filter sizes to capture diverse image features. ResNet-101 and ResNet-50 introduce residual connections to alleviate the vanishing gradient problem and facilitate training of deeper networks. These models offer varying trade-offs between accuracy, computational complexity, and memory footprint, making them suitable for different deployment scenarios.

ResNet-50 Model:

ResNet-50 stands as a pinnacle in the landscape of convolutional neural networks (CNNs), revered for its robust architecture and unparalleled performance in various computer vision tasks. Comprising 50 layers, this formidable model represents a culmination of advancements in deep learning research, particularly in addressing the challenges associated with training exceptionally deep networks.

At the heart of ResNet-50 lies the concept of residual learning, epitomized by residual blocks equipped with skip connections. Unlike traditional architectures where each layer directly feeds into the subsequent layer, ResNet-50 introduces residual connections that bypass several layers. This ingenious design alleviates the vanishing gradient problem, a common hurdle encountered in training deep networks, by facilitating smoother gradient flow and more efficient optimization.

By enabling the training of exceedingly deep networks, ResNet-50 empowers practitioners to capture intricate hierarchical features from input images, leading to superior performance in tasks such as image classification, object detection, and feature extraction. Its ability to learn rich representations of visual data has made it a cornerstone in the development of state-of-the-art computer vision systems, contributing to advancements in fields ranging from autonomous driving to medical imaging.

Furthermore, ResNet-50's efficiency and effectiveness have transcended conventional boundaries, extending its utility beyond traditional domains. From powering real-time object recognition systems in surveillance applications to facilitating image-based diagnostics in healthcare, ResNet-50 continues to leave an indelible mark on the forefront of technological innovation.

ResNet-101 Model:

ResNet-101 emerges as a formidable extension of the ResNet architecture, heralding a new era of deep learning prowess with its 101 layers. Building upon the foundation laid by its predecessor, ResNet-50, this model embodies the relentless pursuit of innovation in convolutional neural network (CNN) design.

At its core, ResNet-101 inherits the groundbreaking concept of residual learning, a paradigm-shifting approach that revolutionized the training of deep neural networks. By leveraging residual connections, ResNet-101 effectively addresses the vanishing gradient problem, a longstanding challenge hindering the training of deep architectures. These skip connections facilitate seamless gradient flow, enabling efficient optimization even across hundreds of layers.

The increased depth of ResNet-101 empowers it to capture richer and more nuanced representations of visual data, unlocking new frontiers in computer vision research. Its versatility is showcased across a myriad of tasks, from the foundational task of image classification to the intricate domains of object detection and semantic segmentation. In image classification, ResNet-101's deep architecture enables it to discern subtle patterns and intricate features, resulting in heightened accuracy and discriminative power. In object detection, its robust feature extraction capabilities facilitate precise localization and recognition of objects amidst cluttered backgrounds. Moreover, in semantic segmentation, ResNet-101's hierarchical representations enable pixel-level classification, paving the way for detailed scene understanding and analysis.

DenseNet-121 Model:

DenseNet-121 emerges as a pivotal variant within the landscape of the DenseNet architecture, offering a profound synthesis of depth and efficiency with its 121 layers. Rooted in the principles of dense connectivity, this model represents a paradigm shift in convolutional neural network (CNN) design, redefining the notion of feature reuse and parameter efficiency.

At its essence, DenseNet-121 embraces a revolutionary architectural paradigm wherein each layer receives direct input from all preceding layers. This dense connectivity fosters maximal information flow throughout the network, facilitating holistic feature extraction and propagation. By promoting

seamless information exchange between layers, DenseNet-121 transcends traditional architectures, mitigating issues of information degradation and fostering deeper insights into the underlying structure of visual data.

Like its counterpart DenseNet-201, DenseNet-121 harnesses the power of densely connected blocks to amplify feature reuse and parameter efficiency. These blocks serve as the building blocks of the network, facilitating comprehensive feature aggregation and propagation across successive layers. The result is a compact yet potent architecture that excels in extracting intricate patterns and nuanced representations from input images.

MobileNet Model:

MobileNet emerges as a pioneering convolutional neural network (CNN) architecture tailored explicitly for mobile and embedded devices, embodying a paradigm shift towards efficient and lightweight deep learning models. With its innovative design and emphasis on computational efficiency, MobileNet revolutionizes the landscape of deep learning deployment in resource-constrained environments.

At its core, MobileNet leverages depth-wise separable convolutions, a novel architectural component that decouples spatial and channel-wise convolutions, significantly reducing computational complexity while retaining expressive power. This separation allows MobileNet to achieve remarkable efficiency gains by performing convolutions in two distinct steps, drastically reducing the number of parameters and computations required compared to traditional convolutional layers.

The hallmark of MobileNet lies in its efficient utilization of parameters and low memory footprint, attributes that are paramount in scenarios where computational resources are scarce or constrained. By striking an optimal balance between model size and performance, MobileNet offers a compelling solution for a wide range of applications, from real-time image classification on mobile devices to edge computing in IoT (Internet of Things) deployments.

Inception V3 Model:

Inception V3 stands as a pinnacle in convolutional neural network (CNN) architecture, distinguished by its innovative inception modules and pioneering approach to feature extraction. With its meticulously crafted design, Inception V3 embodies a fusion of depth, versatility, and computational efficiency, making it a cornerstone in the realm of computer vision.

At the heart of Inception V3 lie its inception modules, architectural components designed to capture diverse image features through convolutional layers with multiple filter sizes. This multi-scale approach enables Inception V3 to extract a rich spectrum of visual information, ranging from fine-grained details to high-level abstractions, thereby enhancing its ability to discern complex patterns and structures within images.

Moreover, Inception V3 incorporates factorized convolutions and parallel branches, architectural innovations aimed at managing computational complexity while maximizing feature representation. Factorized convolutions decompose traditional convolutions into smaller, more manageable operations, reducing the overall computational burden without sacrificing expressive power. Parallel branches further augment feature extraction capabilities by facilitating simultaneous processing of input data through multiple pathways, enabling holistic analysis and interpretation of visual stimuli.

VGG16 Model:

VGG16 stands as an iconic figure in the realm of convolutional neural network (CNN) architectures, revered for its simplicity, elegance, and remarkable effectiveness in diverse computer vision tasks. With its deep stack of convolutional layers and meticulously designed architecture, VGG16 represents a cornerstone in the evolution of deep learning models.

At its core, VGG16 embodies a straightforward yet powerful design philosophy, characterized by its multiple convolutional layers equipped with small 3x3 filters. These compact filters, combined with the application of max-pooling layers for spatial down-sampling, enable VGG16 to systematically extract hierarchical features from input images. By cascading these layers, VGG16 can progressively learn increasingly abstract and complex representations of visual data, thereby empowering it to discern intricate patterns and structures within images.

One of the defining features of VGG16 is its deep architecture, comprising a total of 16 layers, including convolutional and pooling layers. This depth plays a pivotal role in the model's ability to learn rich and discriminative features from input images, enabling it to achieve state-of-the-art performance across a wide range of computer vision tasks.

The versatility of VGG16 is evident in its widespread adoption across various domains of computer vision. From image classification to object detection and segmentation, VGG16 consistently demonstrates its prowess in accurately categorizing and understanding visual data. Its robust performance and ease of implementation have made it a go-to choice for researchers and practitioners seeking reliable solutions to complex vision problems.

Furthermore, VGG16's architectural simplicity and effectiveness have transcended traditional boundaries, inspiring countless iterations and adaptations in the deep learning community. Its influence can be seen in subsequent architectures and methodologies, shaping the trajectory of research in computer vision and deep learning.

4. DATASETS

The terrain classification dataset represents a comprehensive collection of approximately 30,000 images meticulously categorized into distinct terrain types: grassy, marshy, rocky, and sandy. This dataset serves as a foundational resource for training and evaluating machine learning models tasked with the challenge of classifying terrain types based on visual cues.

Structured with meticulous care, the dataset is divided into three subsets: training, testing, and validation sets. This partitioning ensures robust model development and performance assessment by providing separate sets for model training, evaluation, and validation. The training set, comprising a substantial portion of the dataset, serves as the bedrock for model training, enabling the algorithm to learn the intricate nuances and characteristics associated with each terrain type. Meanwhile, the testing set offers an independent assessment of the model's performance, allowing researchers to gauge its generalization capabilities and identify potential areas for improvement. Lastly, the validation set acts as a final litmus test, validating the model's performance on unseen data and ensuring its reliability and accuracy in real-world scenarios.

Each image within the dataset is meticulously labeled with its corresponding terrain type, providing ground truth annotations essential for supervised learning. This labeling scheme not only facilitates model training but also enables rigorous evaluation of model performance against predefined ground truth labels. As a result, researchers and practitioners can confidently develop and test terrain classification algorithms, leveraging the rich and diverse dataset to drive advancements in machine learning and computer vision.

5. RESULTS

In our analysis of the Kaggle terrain classification dataset, we implemented six different models: ResNet50, ResNet101, MobileNet, InceptionV3, DenseNet121, and VGG16. After rigorous evaluation, we found that MobileNet achieved the highest accuracy of 90%, followed by DenseNet121 with 85% accuracy.

The success of MobileNet in achieving the highest accuracy highlights its effectiveness in balancing model complexity and performance, making it well-suited for tasks with limited computational resources or strict deployment requirements.

DenseNet121 also performed commendably with an accuracy of 85%, showcasing its ability to capture intricate patterns within the dataset despite its deeper architecture. This emphasizes the importance of leveraging models with varying complexities to explore the trade-offs between performance and computational efficiency.

Moreover, these results underscore the significance of choosing appropriate model architectures tailored to the characteristics of the dataset and the computational constraints of the task at hand. Experimenting with a diverse range of models allows for a comprehensive understanding of their capabilities and limitations in solving specific classification problems like terrain classification.

Moving forward, further analysis could delve into understanding the specific strengths and weaknesses of each model in classifying different terrain types, potentially leading to insights that can enhance model performance or inform the development of specialized architectures optimized for terrain classification tasks.

6. CONCLUSION

Our journey through terrain classification using image datasets has been both illuminating and fruitful, culminating in the successful deployment of various transfer learning models and the attainment of a commendable accuracy rate of 90% with MobileNet. This achievement serves as a testament to the efficacy of harnessing pre-trained models for such intricate tasks in computer vision.

As we reflect on our accomplishments, we are mindful of the evolving landscape of terrain analysis and the inherent challenges that lie ahead. Looking forward, our aspirations extend beyond mere performance metrics; we are driven by a steadfast commitment to innovation and advancement in the field.

One of our primary objectives for the future involves the expansion of our dataset to encompass a more comprehensive array of samples, including diverse terrains and environmental conditions. By broadening the scope of our dataset, we aim to fortify the robustness and generalization capabilities of our models, thereby ensuring their efficacy across a broader spectrum of real-world scenarios.

Moreover, our roadmap for future research endeavors includes the exploration of advanced architectural paradigms such as vision transformers and encoders. These cutting-edge approaches hold immense potential for revolutionizing terrain classification by offering novel mechanisms for feature extraction and representation. By embracing these innovative techniques, we endeavor to transcend existing limitations and push the boundaries of performance in terrain analysis tasks.

7. REFERENCES

- [1] Yang, J., Xu, J., Lv, Y., Zhou, C., Zhu, Y., & Cheng, W. (2023). Deep learning-based automated terrain classification using high-resolution DEM data. *International Journal of Applied Earth Observation and Geoinformation*, 118, 103249.
- [2] Bai, C., Guo, J., & Zheng, H. (2019). Three-dimensional vibration-based terrain classification for mobile robots. *IEEE Access*, 7, 63485-63492.
- [3] Shi, W., Li, Z., Lv, W., Wu, Y., Chang, J., & Li, X. (2020). Laplacian support vector machine for vibration-based robotic terrain classification. *Electronics*, 9(3), 513.
- [4] Christie, J., & Kottege, N. (2016, May). Acoustics based terrain classification for legged robots. In 2016 IEEE international conference on robotics and automation (ICRA) (pp. 3596-3603). IEEE.

-
- [5] Pingel, T. J., Clarke, K. C., & McBride, W. A. (2013). An improved simple morphological filter for the terrain classification of airborne LIDAR data. *ISPRS journal of photogrammetry and remote sensing*, 77, 21-30.
- [6] Li, S., Xiong, L., Tang, G., & Strobl, J. (2020). Deep learning-based approach for landform classification from integrated data sources of digital elevation model and imagery. *Geomorphology*, 354, 107045.
- [7] Zhang, L., Chen, Z., Zou, B., & Gao, Y. (2018, July). Polarimetric SAR terrain classification using 3D convolutional neural network. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium* (pp. 4551-4554). IEEE.
- [8] Yu, Z. (2020, September). Research on Remote Sensing Image Terrain Classification Algorithm Based on Improved KNN. In *2020 IEEE 3rd International Conference on Information Systems and Computer Aided Education (ICISCAE)* (pp. 569-573). IEEE.
- [9] Zürn, J., Burgard, W., & Valada, A. (2020). Self-supervised visual terrain classification from unsupervised acoustic feature learning. *IEEE Transactions on Robotics*, 37(2), 466-481.
- [10] Abete, J. M., Abooe, A., Acharya, P., Afzal, M. R., Agnus, J., Ahlers, D., ... & Cai, C. J. (2021). 2021 Index *IEEE/ASME Transactions on Mechatronics* Vol. 26. *IEEE/ASME TRANSACTIONS ON MECHATRONICS*, 26(6).